

### **REMARKS**

The Examiner rejected claims 1-20. Particularly, the Examiner rejected claims 1-20 under 35 U.S.C. §102(e) as being anticipated by U.S. Pat. No. 6,298,451 of Lin ("Lin"); and rejected claims 1-20 under §102(e) as being unpatentable over U.S. Pat. No. 6,564,131 of Hickman et al. ("Hickman"). Applicant has amended independent claims 1 and 8 to include elements of certain dependent claims that are not disclosed by any of the prior art. Particularly, claim 1 was amended to include storage server, metadata, and gateway service elements of claims 4, 5 and 6, respectively; and claim 8 was amended to include storage server and metadata service elements of claims 12 and 10, respectively. Claims 2-4, 9, 10, 12 and 15-20 have been canceled. None of the prior art discloses or suggests a scalable storage system that can handle requests that can affect multiple back-end servers and that includes separate services for metadata and storage, as set forth in Applicant's pending claims. The Examiner's early reconsideration and withdrawal of the rejections are respectfully requested.

#### *Lin*

Independent claims 1 and 8 have been rejected as being anticipated by Lin. The independent claims have been amended to include elements that are neither taught nor inherent in Lin. Particularly, claim 1 has been amended to include separate storage server, metadata, and gateway service elements of claims 4, 5 and 6, respectively; and claim 8 was amended to include separate storage server and metadata service elements of claims 12 and 10, respectively. Lin does not teach a storage system with separate metadata and storage server elements, nor a system that can manage requests that may affect multiple back-end servers.

Lin describes a load-balancing distributed server scheme, where Lin's "gateways" act only as forwarding agents, and where any given request is handled by one back-end server. A given "class" of request may be handled by several back-end servers, and a directory service records which back-end servers handle each class of

request. When a new request arrives at a gateway, the gateway asks one of the directory servers for a reference to a back-end server handling that class of request, and then forwards the request to a specified server. Each request is handled by one specific server. Furthermore, Lin does not disclose a "storage system" with separate components, such as a metadata service and a storage service.

Claim 1, as amended, recites a storage system with separate metadata, gateway and storage services. The claimed invention provides a system that can distribute the functions involved in handling a given storage system request. Further, in the claimed invention, a given request may affect *multiple* back-end servers, unlike the case in Lin. Indeed, most requests for file service affect in principle multiple back-end servers in a scalable system such as Applicants, assuming that some handle metadata for a given file and others handle file data for the file, and assuming that the file hierarchy is distributed over multiple metadata servers.

For example, a rename operation, where the old and new parent directories, the file being renamed, and a preexisting second file having the new name all reside on different metadata servers, will require the participation of four separate back-end servers. Lin makes no provision for coordinating operations involving multiple back-end servers for a single request. Rather, Lin teaches assigning a task to a single back-end server.

Lin is essentially a load balancing scheme for back-end instances that are equivalent, not an integrated scalable storage system. That is, Lin applies only to completely parallel scalability (typically read-only) or to cases where updates are partitioned and an update to a given partition is independent of updates to other partitions. The present application is concerned with cases where operations may generally be applied to multiple partitions (although need not be applied to multiple partitions in every case), and further accounts for gateways handling many requests (particularly for retrievals) themselves, rather than being simple forwarding agents.

Lin does not describe this approach to providing a file storage system. Lin only describes having various separate services co-hosted on a variety of back-end servers, with no discussion of how this might apply to file storage in a generally scalable way (as opposed to the trivial case of having a large number of unrelated small file systems, each on a different service). The present application supports scaling a single file system, not the mere aggregation of independent services.

Furthermore, Lin does not disclose or suggest a system having an independent metadata service. The Examiner interprets "metadata service" to be a "service defining task and locating server to perform task." This interpretation contradicts the commonly known definition of "metadata" and the description of the metadata service set forth in the pending application. In the pending application, Applicants explain that the metadata service "can access metadata for various files in a storage system" (page 6), and that "metadata" is commonly known to include "predetermined information on files contained in the storage system" (page 14). The information may include information regarding the hierarchy of the file system and the location of the files (Pending application at page 14, lines 14-21). The interpretation proposed by the Examiner is not related to metadata or predetermined information on files contained in the storage system, such as file location and file system hierarchy information. The entire disclosure of Lin does not contain one reference to "metadata" or to a separate service that provides access to "predetermined information on files contained in a storage system." Thus, Lin fails to teach the metadata service element of the claimed invention.

For all of these reasons, Lin does not teach all of the elements of claim 1. Therefore, Lin cannot anticipate claim 1 or any claim depending from claim 1 (e.g., claims 6 and 7). Allowance of these claims is respectfully requested.

Like claim 1, claim 8 recites a storage system that can handle requests that affect multiple back-end servers, and that includes separate storage service and metadata service components. As set forth above, Lin does not disclose a scalable

storage system that can handle requests that affect multiple back-end servers. Moreover, Lin does not disclose or suggest a storage system with a separate metadata service. For at least these reasons, Lin cannot anticipate claim 8 or any claim depending from claim 8 (e.g., claims 11, 13 and 14).

*Hickman*

Pending independent claims 1 and 8 have been rejected as being anticipated by Hickman (although the Examiner does not explain how Hickman anticipates claim 1, and begins the discussion with claim 8). Since both claims 1 and 8 include elements not found in Hickman, Hickman cannot anticipate either claim.

As in Lin, Hickman fails to disclose a separate metadata service including a plurality of metadata servers. The Examiner cites to two sections of Hickman in support of the assertion that Hickman discloses the claimed metadata service (i.e., col. 5, lines 45-56 and col. 6, line 55 – col. 7, line 21). Neither of these sections discloses the claimed metadata service. In fact, as in Lin, there is not one reference to “metadata” in the entire Hickman patent.

The cited sections of Hickman refer to a “connection manager 140” and a “storage access module 160” having a “partition map cache 172.” The partition map merely records “client-based” partitioning predicates and cannot be construed to provide a metadata service. Hickman’s partitioning is based on the client, meaning that file storage for a given client is not scalable (it is all in one partition) and access to data for a given client is not scalable. Contrary to the Examiner’s assertion, these elements cannot be construed to provide a metadata service. Rather, at most, these elements provide just a form of cluster directory service. A metadata server, as set forth in the pending application, stores information about the file system hierarchy, the identity of files, and the location of the file data. Hickman does not disclose any such server.

Furthermore, the metadata service in the claimed invention is distributed over a plurality of servers that each access metadata independent of the metadata accessed by other servers (claim 1) or that are "functionally de-coupled from one another" (claim 8). The connection manager 140 and storage access module 160 cannot be collectively construed to fit within either definition. That is, for each file access in the Hickman system, the connection manager 140 and storage access module 160 must both operate together to provide the access, and thus, they are not "functionally de-coupled" or "independent" of one another. Specifically, every time a file is accessed the connection manager 140 selects a web server 145 to handle the communication session (col. 5, lines 45-56), and storage access module 160 retrieves client-specific data from the storage system 104 (col. 6, lines 59-63). As a result, Hickman fails to disclose or suggest the claimed metadata service, which requires a plurality of "independent" or "functionally de-coupled" servers, as recited in claims 1 and 8.


For all of these reasons, Hickman does not anticipate claim 1 or 8 or any claim depending from claim 1 or 8 (e.g., claims 6, 7 and claims 11, 13, 14).

### **CONCLUSIONS**

For all of these reasons, Applicant respectfully asserts that all pending claims 1, 6, 7, 8, 11, 13 and 14 are in condition for allowance. The Examiner's early reconsideration is respectfully requested. If the Examiner has any questions, the Examiner is invited to contact Applicant's attorney at the following address or telephone number:

David Alberti  
c/o Patent Department  
GRAY CARY WARE & FREIDENRICH LLP  
2000 University Avenue  
East Palo Alto, CA 94303-2248

Respectfully submitted,  
Gray Cary Ware & Freidenrich LLP

A handwritten signature in black ink, appearing to read 'D. Alberti', written over a horizontal line.

David Alberti  
Reg. No. 43,465

Dated: June 2, 2004